

SAM (REX/Ops) Operations Meeting 2008.09.23

Adam (typist), SteveW, SteveS, Robert, Art, Eric, Angela, Rick St. D.

Meeting duration 40 minutes.

For reports submitted to the meeting...

The report text is in this font

Comments made during the meeting are in this font inline with the report

Releases:

No scheduled releases.

Here are the unscheduled tickets...

T	Key	Status	Summary	Assignee	Reporter	Res	Cr
	SAMDEV-5	 Open	Improve project robustness against DB failures	Adam Lyon	Robert Illingworth	UNRESOLVED	02
	SAMDEV-1	 Open	Convert station to use omniORB	Robert Illingworth	Robert Illingworth	UNRESOLVED	14

Note: If the idle project killer (which CDF has set for 2 weeks) kills a project while the DB is down, the station will continue onward as if the project is done, but the project is not marked as complete in the DB. So when a file for that project gets uncached, the uncaching fails. Better to have the idle project killer wait forever until the DB is back. Make a subticket for SAMDEV-5.

D0 is running DB server 8.5.3.

CDF is running DB server 8.4.5. (8.5.3 is not in integration yet -- 8.5.3 had the rollback protection. Angela will put in integration).

DØ:

** There was no SAM shifter during Sept. 19-21, 2008.*

** Sept. 16/17 : Mike reported "All jobs going through samgfw02 failed" at night (Sept. 16). Robert on*

Sept. 17 said it seemed a transient problem due to samgfw02 returning to service unexpectedly.

samgfw02 was turned on unexpectedly, so we had two forwarding node #2's at the

same time. Once noticed we turned off d0srv015.

* Sept. 19: (DZSAM-266) : A user complained about the slowness in cabsrv2. Robert found :*"Thanks to an unfortunate interaction between the power outage, the satabeast RAID arrays, and the way SAM works, a sizeable chunk of the cached files on cabsrv2 were deleted yesterday. The slowness is because the SAM station is having to read most of your files from tape. And the errors are because one of the tapes is temporarily unavailable."*

SATAbests were very slow to start up after the planned power outage - hosts turned on first quickly with empty mount points. SAM station noticed this and removed the files from the DB. Lost many TB of files. Solution was to make the mount point write protected so that the station cannot make a boo directory when the real disk is not mounted. Robert did this for all of the servers that were not mounted. It was only a specific generation of SATAbest that had this problem.

Unresolved issues:

DZSAM-267: Mike reported that "Condor expression-evaluation script used to take ~40 minutes but it took 2.5 hours Sunday evening".

Will look at log files.

DZSAM-262/263/264: KCA CA updates.

Parag sent out that the CA cert for KCA expires on 10/8. Forwarding and queueing nodes and cache nodes and all offsite nodes need to be updated and durable locations. Including OSG and LCG -- all SAMGrid nodes need this upgrade. Instructions are in the ticket.

DZSAM-256 : transfer issues between d0srv071 and Imperial College SRM --- confirming with Willem whether this ticket may be closed.

Since we got transfers to work, Closed - DZSAM-54 is still open (the main ticket).

Yesterday - all of the server processes were down on the Lancaster forwarding node.


DZSAM-252: User Smain has confirmed that instructions are sought for how to create metadata to store files.

Still need to talk to him.

DZSAM-54: This has been worked on actively and Willem has just reported another problem.

Production Efficiency:

Monday 2008 September 15 - Monday 2008 September 22



Good running this week.

1. **Thursday:** Scheduled downtime. Memory on forwarding node #2 replaced.

Some nodes did not turn on after the Thursday shutdown.

CDF:

1. CDFDH-258/ Help desk ticket 121711 open
Angela Bellavance updated CDFDH-258: Ticket #: 121711

r-fcdfdata325-1

Status: Waiting (was: Open)

Machine is awaiting hardware replacement.

2. CDFDH-257/259

=====

r-fcdfcache01-1 to 4 offline pings ok

r-fcdfdata318-1 to 4 offline pings ok

The fcdcache machines are in the process of being brought up,
so you can ignore them for now.

fcdfddata318 just got moved and probably the dCache services need
to be restarted. I'll check. Angela

pool 318 is alive now.

3. LT03_01 hardware problem: still a problem

=====

Thanks for your note. LTO3_01 has been down

since Friday because of a hardware issue. A service call has been placed.

4. Lumi Scripts

User running Lumi script asked about getting LUMI on top ntuple.
There are a variety of these lumi scripts and the best is one that Randolph put together. We can mail it 'round but would be best if supported somewhere. Also, these scripts can be hard on the offline db if not used properly.

5. Downtime and recovery

fcdfora4 was back at 11:28
Enstore was back at 13:20
dCache was available at 15:48
CAF was up at 18:00
CDF Web pages had a lot of items to be fixed up.

Noted lots of various other problems. We had to restart two csl fss's (not sure if related to downtime).

Server of pnfs database failed and required attention.

There was a rapidly growing demand for tapes but that dropped off quickly.

7. Tape in drop-off

NOTALLOWED 0919-0940 none) CDF-LTO4F1
cdf.bpel0m.cpio_odc
Tape in drop-off in robot - tech coming 9-19-08

8. Misc.

=====

Some filesize mismatches

Mail from "herber" instead of a common account

Diskpool nodes repaired.

The diskpool file servers fcdfdpool01 and fcdfdpool02 have been successfully transitioned to new GCC housing

Restarted SAM TV once

One stuck job cdf_oper. What does this mean?

Eric says these are calibration jobs - talk to Willis

A job from Japan was pounding on dcache - Stefan Lammel talked to the user and that was fixed.
Since Friday because of a hardware issue. A service call has been placed.

Had this error 09/18/08 18:54:29 to 09/18/08 21:14:30 every 10 min then it went away:
~~4. Lumi Scripts~~

09/18/08 21:14:30 cdf-fncdfsrv0.SM.FileUncachedInDBTimer 13958: Cannot uncache file in DB:

CORBA::UserException caught in FileUncachedInDBTimer::dolt()

See the discussion above under Releases.

MINOS:

Kreymer will be on vacation 26 Sep through 12 Oct.

See you Oct 14.

Will be away for next two meetings.

Enstore -

DCache -

SAM -

Oracle -

OPEN ISSUES

IT 1146: cannot test existence of sam tape location

IT 2071: quadratic project startup time growth

PLAN -

DB:

D0 :

Alert(s)

- Added 6GB to OPERATIONAL_IDX01

CDF:

- OS and DB patches on cdfopr2 done - Sep 18. Also increased the sessions parameter to 900

Diana wants to install the Oracle Security patch to cdfstrm1 (the offline replica database). When cdfstrm1 is brought down users will fail over to cdfopr2 - will look for a quiet time for cdfopr2 to do this.

Alert(S):

MINOS:

Alert(s):

Enstore/dCache:

Enstore Operations report

- Release Report:

+ Scheduled Downtimes:

CDFen downtime: Thursday 9/18 took 1 hour longer than anticipated. Otherwise it went well.

- Change Report:

+ Tentative downtime on Tuesday Oct 7, 2008 for D0EN.

+ Tentative downtime on Thursday Oct 16, 2008 for STKEN.

+ Tentative downtime on Wednesday Oct 8, 2008 for CDFEN.

- Problem Report:

+ Unscheduled Downtimes:

- cdfensrvln began to reboot intermittently on it's own. The unit was

replaced on Friday 9/19/2008 the day after the CDFEN downtime.

We

believe this to be fallout from the downtime. There was problem with one of the virtual network interfaces that failed to come

up

at boot and was not properly monitored. This extended the downtime

of the server. 5 hours down.

- SL8500#3 had a BOT fail on Wednesday 9/17/2008. Sun replaced it.

CDF and D0 experienced 4 hours of no availability of this library.

- SL8500#3 Placed a tape in the dropoff. The library does this with a

tape if it loses track of it's real location. We had to pause the

D0 and CDF LTO4 Library Managers so that we could open the Library

and retrieve the tape from the dropoff. CDF and D0 experienced

Alert(S):

minutes of no availability of the SL8500#3.

- stkendca9a became unresponsive. Node was rebooted. Node logs are being investigated.

MINOS:

- Capacity Report:

+ Agilysys has been awarded the the SL8500 purchase of 2 SL8500's.

We

are having a pre-install meeting Monday 9/22, 2008. Sun wants to deliver the hardware Tuesday, tomorrow.

- Incident Report:

Enstore/Cache incidents reported last week.

Enstore/Cache:

	7/14	7/21	7/28	8/4	8/11	8/18	8/25	9/8	9/15	9/22	
17 incidents	21	48	5	25	4	36	19	25	26		enstore mover
20 incidents	34	50	30	24	6	0	48	20	46		tape incidents
0 incidents	0	1	1	0	1	0	2	0	0		tape drive hw
1 incidents	1	1	0	2	1	3	0	0	3		library hw
0 incidents	0	0	5	1	0	0	0	7	0		enstore mover hw
0 incidents	0	0	2	1	0	1	1	0	0		enstore server hw
2 incidents	0	0	0	0	0	1	3	4	1		dcache pool hw
0 incidents	0	0	0	0	0	0	0	0	0		dcache server hw
0 incidents	0	0	0	0	0	0	0	0	1		pnfs hw incidents
0 incidents	0	0	0	0	0	0	0	0	0		pnfs sw incidents
3 incidents	3	0	3	0	1	0	0	0	0		dcache sw incidents
0 incidents	1	0	2	0	1	1	0	1	0		enstore sw
0 incidents	0	0	0	0	1	0	0	0	0		enstore config
1 incidents	0	0	0	0	0	0	0	3	0		dCache config

7 outage hours	4	0	0	2	0	12	0	0	9.5		unsched. enstore
0 outage hours	0	0	0	0	0	0	0	0	1		unsched. dCache

Data Integrity:

- VSN PRQ585. Still waiting for Sun to return tape.
- Enstore is investigating an issue with tapes not being filled to capacity. This appears to be LT04-related only at this point. Here are plots of LT04 tapes plotted time sorted by filled tapes. The dates start in Mar, 08 and end today. The part of the plot that shows tapes being filled to normal capacity is after the SDE upgrade

of Stken on July 24, 2008. The D0 tapes that show up in the plot as

not full were written in the GCC SL8500 during Duplication/migration.

<http://www-ccf.fnal.gov/STK800TimeSortPlot.jpg> for Stken (SL8500#2).

and

<http://www-ccf.fnal.gov/D0800TimeSortPlot.jpg> for D0en
(SL8500#3).

We have tested reading some of these tapes and have had no problems

reading them. We are still investigating how this could happen, and

why after the upgrade to SDE this issue went away. Our theories revolve around tape drive F/W update and mover code.

- Migration work flow status: These plots are being fixed this week,

they will reflect better the migration status.

LTO1 Migration: 0 / ~1200

LTO2 Migration: 2853 / ~8500

9940A Migration: 1728 / ~4500

9940B Migration: 0 / ~24000

LTO1 Migration: 0 / ~1200

- SSA group is interviewing for a Full time position. Interviews and

phone screens in progress.

FNDCA:

We have several bugzilla tickets opened by MINOS:

Details in the "additional information" section of this report.

CDFDCA:

A number of pools have been configured to auto-restart on reboot
09-19 - restarted PnfsManager which has been hanging since evening

09-18. cdfensrvln developed h/w problem and has been swapped.

During

repair and shortly after, a file could not be read by tape.

Normal

operations resumed at about 18:20 09/19

New pool nodes fcdfcache01-09 "officially" added to configuration after power outage (we were waiting for end of shutdown).

We restarted all dCache after power outage in gcc when SSA worked on

pnfs node.

After shutdown, the pnfs server node rebooted several times. At least

on a few occasions the pnfs server did not start properly on reboot.

What about the DNS problem that caused enstore to stop transferring files? The config files with secondary and tertiary DNS servers listed were there. Speculation is that <http://www-ccf.fnal.gov/STK800TimeSortPlot.jpg> for Stken primary DNS server was working enough to prevent the fail over but not enough to actually work: gi and

primary DNS server was working enough to prevent the fail over but not enough to